



Artificial Intelligence: A Medium that Hides Its Nature

Avon Huxor

The field of artificial intelligence (AI) is unusual in that it has an origin story. In ‘A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence’, the founders, back in 1955, argued that every feature of intelligence can be so well defined as to be amenable to simulation (McCarthy et al. 2006). This proposal led to an (extended) event that took place the following year (Moor 2006). The proposal document is believed to be the first time that the term *artificial intelligence* was used. The goal of the field is explicit in this name. That is, the aim of AI is to simulate intelligence to such a level that it becomes indistinguishable from *human intelligence* itself. I refer to this view as the traditional one.

In this essay I argue the case for a view in which AI can be better seen as a medium—a smart medium—with more in common with writing

A. Huxor (✉)

University of Exeter, Exeter, UK

e-mail: a.huxor@exeter.ac.uk

and cinema than with the thinking machines of science fiction. This approach delivers a number of benefits over the traditional approach. Its origins still hold considerable sway over AI and have led to a tendency to ignore alternative perspectives on what the technology can do. But, I argue, we can use the ideas that have been developed over decades to understand existing media, to throw new light on smart media, on their authorship and reception. More importantly, we can design new AI applications with a media approach built-in from conception to deployment. Given the impact of printed books, cinema and television on science, art and culture, I cannot imagine a greater ambition than using the powers of this new technology to expand human expressive capabilities—rather than attempting to imitate and replace them.

First, however, it is useful to try and understand some of the reasons that the traditional view of AI is so resilient.

Dialogue and Anthropomorphism

At the time of its foundation, in the 1950s, AI found a ready-made metric against which many of its ambitions could be tested: the Turing Test. This test (Turing 1950) is well-known, so I will not describe it in detail. Suffice to say that it is conventionally understood as setting up conversations, over computer terminals, between a questioner and two other participants: one of which is a computer and the other is human. Through a dialogue involving typing text—only questions and answers—through the terminals, the questioner must decide which of the other two subjects is a human and which is a computer. The test is passed by the computer when the questioner cannot distinguish its responses from the human player.

Many issues have been identified with the Turing Test as a measure of machine intelligence, however, which has given rise to rich literature on the topic(see, e.g., the review of Saygin et al. (2000)). There is one that I would like to address, and it revolves around a tendency for humans to identify agency and intent in the world around us. This urge can be seen

in the classic study of Heider and Simmel who created a film¹ in which very simple cartoon circles and triangle moved in and around a simple box shape, yet viewers attributed intent and narrative to these simple shapes(Heider and Simmel 1944). The intent, of course, can actually be found in the mind of the animator. More recent empirical studies have found that humans unconsciously apply the same social heuristics used in human interactions to computers because they bring to mind similar social attributes as we have with our fellow humans(Reeves and Nass 1996).

The dialogue style of interaction used in the Turing Test is a strong trigger for such attribution of agency to a machine: “It can be shown that attributing to an artifact the position of interlocutor in a dialogue implies dealing with it as if it were endowed of the features characterizing human mind...”(Airenti 2015: 117). A clear example can be found in the reaction of users to ELIZA(Weizenbaum 1966), one of the earliest AI systems, one of whose variants mimicked the interaction with psychotherapist. It was very simple in its functionality: for the most part asking a few standard questions and then mirroring back much of the human user’s answers. For example, it might take a word (X) in the answer and ask in return “Tell me more about...X”. Its creator was horrified that users, even those who knew about its simple nature, would become drawn into conversation with ELIZA. He feared that such technologies blurred the line between the human and the machine. Despite the trivial nature of ELIZA’s actual programming, the dialogue was the key feature of ELIZA that evoked a human reaction to anthropomorphise the machine.

Dialogues can also be found in other technologies. As a student in the late 1970s I recall that some friends would find it amusing to play an informal game. They would leave an answer-phone message (this took place in the days of a message being recorded on a small tape machine that was connected to the telephone) and tried to see how long they could confuse the caller into thinking they were actually answering. They would record phrases such as “Oh, hi there...” and “what’s up then?”

¹ <https://blogs.scientificamerican.com/thoughtful-animal/animating-anthropomorphism-giving-minds-to-geometric-shapes-video/>.

with suitable pauses that invite the caller to reply. One could listen to the dialogue between the caller and the tape machine on the recording. This seemed to work better than one might expect. It would take some time before the caller realised that they were talking to a recording rather than a real person. The conversational effect led to the caller interpreting the recorded responses as live speech. People would even allow for talking over and non-sequiturs as these are not unexpected in a telephone call.

The ELIZA effect, in which the human urge to find agency in the world around us, given minimal cues, should make us cautious of taking the successes of many AI systems at face value. Throughout the history of AI, however, there has been a parallel commentary, which takes an alternative perspective. This is a view that does not see the computer as another *agent*, but as a *medium* through which we can communicate with other people, those involved in the creation of its content. This idea is one that has as long history too, but one that has—thus far—always been a quiet voice compared the dominant and often booming narrative of traditional AI.

Half a century ago, one critique of the Turing Test specifically suggested that the computer used in the Test is really a medium, through which human ideas are communicated, between the computer programmer and the human interrogating the system: “It is really a battle of wits between questioner and programmer...Whatever thinking activity is going on is plainly in the brain of the programmer, not in the mechanism by which his answers are transmitted” (Purtill 1971: 291). Similarly, Paulson noted that the Test can be directly compared to the creation of fictional characters in literature: “The Turing Test requires intelligent or person like behavior to be manifested strictly in written language. The problem is in effect an interactive version of the kind of situation we encounter routinely in reading literature, most obviously (though not exclusively) with respect to effects of character in realistic novels” (1989: 299). It is not surprising that the value of discussing alternative approaches to AI gained some traction during an AI winter, where AI—having fallen short of its ambitions—found itself on the back foot (Hendler 2008). Paulson wrote during one such AI winter.

At about this same time, other scholars generalised this view beyond the Turing Test and argued that we can see all AI systems as a ‘medium’—one that allows humans to externalise and communicate their intelligence to others—rather than being autonomous intelligent agents, and that writing as a medium provides particularly valuable lesson. Bolter (1991: 188) has summarised much of the argument of this essay, “books do not replace human intelligence; they present human intelligence in another form. Computers should be understood in the same way”. Likewise, Sokolowski (1988: 51) has written that it would be more fruitful if “instead of simply comparing computers and brains, we should also compare the ‘reasoning’ of the machine with the ‘stating’ of the sign, and examine storage and representation as they occur in the machine and in writing”. These ideas, however, were not taken up by the AI community. There are likely a number of reasons for the lack of interest. These authors are outsiders—not from the AI community and they did not publish in AI journals or attend AI conferences; the authors I cite above all come from a background in the humanities. However, it is also very likely that the new metaphor of AI as a medium simply strayed too far from the stated aims of the AI enthusiasts. It was around this time, in the late 1980s and early 1990s, that I too became convinced of the need for a new approach to AI (Huxor 1994) and experienced the reluctance of many in the field to move away from its traditional goals.

However, there have been rumblings—even within the ranks of AI—against the tendency for people working in the field to succumb to wishful thinking in how they describe the technology. McDermott (1976) criticised his fellow researchers for this attitude decades ago. For example, he compares the use of terms in two AI programs that use very different terms for the same idea. One uses the term ‘GOAL’, while another prefers ‘FETCH & TRY-NEXT’. As he notes, the word GOAL gives a feeling of power, but one that is an illusion. Lipton and Steinhart (2019) and Bender and Koller (2020) argue that such loose use of terms to imply greater intelligence than really warranted continues to the present day. For example, they cite authors who describe an machine learning model that “understands sentence relationships”. McDermott’s recognition of the power that comes from working within the conventional AI paradigm is instructive. It is interesting to look at those ideas

which AI itself considers to be its historical and cultural context, because they give us an indication as to the motivations that underlie the project. Thus, in one early publication, a panel of many pioneers of AI refer to the automata created by the Greek god Hephaestus and note that “for humans to behave like gods - because godlike it is to imbue the inanimate with animation - is hubris indeed, and no opponent of artificial intelligence has failed to express shock at the blasphemous behavior of humans who aspire to divinity” (McCorduck et al. 1977: 951).

Instead of aspiring to divinity, could we not find a way to use the technology to give to tools to our fellow humans to express themselves more fully and explore new creative and rhetorical possibilities²? Consider the impact that the written and printed word has had upon culture: law, literature and science. It is to writing that we briefly turn.

AI and Writing

Due to the textual nature of the Turing Test and the textual nature of much AI work, we can look at what writing, as a medium of communication, can contribute to our alternative understanding of AI.

In reading a written text, a burden falls on the reader to reconstruct the writer’s attitude to the text. An orator could simply use the tone of voice, for instance, to convey this. To allow writing to do the same required the invention of new forms. For example, writing practices include the use of the question mark to convey that a statement is a question, yet it does not represent the rise in tone at the end of the sentence used in most forms of spoken English. Compare the spoken form of:

The house is on fire.

And

The house is on fire?

² On the topic of rhetorical possibility, see Ullmann’s experiment with AI translators, in this anthology.

This idea can be extended. In reading a written text, an additional burden falls on the reader to reconstruct the writer's *attitude* to the text. An orator can simply use the tone of voice to convey this. To allow writing to do the same required the invention of new forms of expression. As Olson notes, a narrator does not need to say, "he insisted that..."; the orator simply uses an insistent tone of voice. But the writer cannot do this. Thus, new verbs—representing speech acts and mental states—were employed. Olson shows how Old English had a limited palette of speech act verbs: believe, know, mean, say, tell, think and understand. Later, in the fifteenth to seventeenth centuries, terms, such as assume, contradict, criticise, explain, infer, suggest, and others, came to be used.

For David Olson, the history of writing "is the struggle to recover what was lost in simple transcription" (1994: 111). So writing is not just about writing down the words spoken but encoding what was not actually spoken but present in the oral form through the context and manner of speech. This was required as the author cannot be consulted as to how the text would be understood. What is needed was an indication of what the author intended—what they might have said had they been present at the act of reading.

Taking this broader view of writing further, AI technologies provide the ability for an *AI text* to convey what the author "would have said" given the chance, but in a more comprehensive manner than the printed word. It allows the inherent reasoning "written down" by the author(s) in the software to be made present to the reader, who is the user of the system. That is, AI extends the written or printed page one step further. The author creates text which allows the reader to see how he/she might have wanted to text to be intended for a range of alternative questions presented to it. That is, the lines of reasoning captured by the AI program can answer questions that the author expects to be asked. We can ask the *text* to explain itself as the author might have been. We have to understand the term *text* more widely here. It is the representation of the authors' knowledge in the machine that is used to generate the visible words—it includes their reasoning, embodied in dynamic computer code. Being dynamic, these texts can only be fully written and read on a computational surface, much as early cinema required a celluloid surface.

SYMPTOMS	DIAGNOSIS	SELF-CARE
Begin Here		
1. Has your cough begun recently?	Yes Go to Question 12.**	
No		
2. Are you very short of breath, and are you coughing up pink, frothy mucus?	Yes Your symptoms may be from a serious condition called PULMONARY EDEMA (fluid in the lungs).	EMERGENCYGO TO THE NEAREST EMERGENCY ROOM RIGHT AWAY.
No		
3. Does your cough produce clear or pale yellow mucus?	Yes You may have a viral illness such as a COLD or the FLU.	Get plenty of rest, and drink lots of fluids. Try over-the-counter medicines to treat your symptoms.
No		
4. Does your cough produce yellow, tan or green mucus?	Yes You may have an infection of the airways such as CHRONIC BRONCHITIS. If you have a fever with shaking chills and are very ill, you may have a more serious infection such asPNEUMONIA.	See your doctor. He or she can prescribe medicine to relieve your symptoms. Get plenty of rest, and drink lots of fluids. If you smoke, stop smoking.
No		
5. Does the cough come with shortness of breath and	Yes These symptoms may be a sign of ASTHMA, a constriction of the airways.	Asthma can be dangerous and should be diagnosed and treated by

Fig. 1 Flowchart for the diagnosis of a cough

The ability of conventional writing to capture some aspects of thinking predates the computer and relies to a great extent on the fixed, spatial character of the written page. Consider the flowchart for the diagnosis of a cough³ part of which is shown in Fig. 1. The reader follows a set of if-then questions, where a ‘yes’ answer will point to either a later question or give a diagnosis (followed by some suggested self-care). This medical flowchart is not complicated and the information structure, the rules, can be represented in the spatial layout of the page. The structure of that reasoning is captured by the nature of the surface on which the text is written: the physical layout has properties that represent the thinking.

Such a series of if-then rules was central to the *expert systems* (see, e.g., Kastner and Hong 1984) of the AI boom in the 1980s. If the diagnosis system was to become more complicated, it would not be practical to place on printed pages—it would be large and cumbersome. The solution was to represent the structure of the rules not by the spatial

³ <https://drgreenmom.com/cough/>.

property of the printed paper surface (or a screen-based equivalent) but by the *computational surface* of the computer. Again, the idea of the computer somehow being a surface sounds counterintuitive but let us remember that the strip of 35-mm film that captures a cinema film is technologically complex, requiring ancillary writing and reading devices.

Once implemented in a computer it is easy to see how the same type of medical knowledge can be expressed in a conversational interface. The computer ‘asks’ the questions in its database and follows the decision tree to present the reader with a diagnosis and recommendations. This dialogue form of interaction would appear, due to the anthropomorphism discussed above, to imbue the system with agency, while the printed, written and format does not exhibit any such property; we see it as an authored text.

Note that these conversational interfaces, in which a textual dialogue takes place, were common in these early systems due to the simpler nature of the technology available at that time. The simple teletype terminal with no graphical abilities required the system designer to have a dialogue interface, one that will likely trigger the ELIZA effect. For example, an early spell checker from 1980 is described by Peterson (1980) which has a dialogue interface. These days we are used to spell-checkers employing a graphical user interface (GUI), in which a word is highlighted (typically with a red wavy underline) as possibly misspelt. Using the mouse, the user obtains a small pop-up menu listing possible correct spellings, the option to add the word to the dictionary (if it is new) or simply ignoring it. But such an interface was difficult in the early days of computing with their simple terminal screens. Thus in Peterson, the user is given an interface in which a misspelt word is highlighted in the text, but a separate part of the screen below contains a dialogue with the user, asking.

“What to do with misspelled token? (Select one)”

with a list of options to choose. This older interface feels like a simple AI, a chatbot, and the user apparently converses onscreen with a spelling expert. In a modern text editor there is no such feeling.

Finding the Author in Knowledge Medium

The arguments above are of a general nature. The question now is how can we start to see through the ELIZA effect and into AI's nature as a medium? Part of the answer could lie in the observation of Bolter that "the strategy of artificial intelligence is always to disguise the computer as text and make of it instead a persona" (Bolter 1991: 184). Our strategy is now to reverse this, to try and locate and highlight the authors that sit behind the AI text, thereby revealing their true nature. This proved challenging as, too often, the literature does not document the authorship of AI systems well.

Let us begin with AI tools for the generation of text, often called robot journalists in the press and seen as a threat to journalists' jobs (Martin 2019). These newly developed systems extend naturally from the discussion above on tabular flow charts and spell-checkers. Press reports for these commercial systems invariably have a stock image of a robot typing at a keyboard, evoking the traditional AI goal.

It is difficult now to get details of how these systems function, but an early description of the work of one company 'Narrative Sciences' is revealing (Levy 2012). This early system took raw data and generated newspaper stories from these data. One example was on the subject of American college sports, where normally it would not be economic to have journalists working full-time writing stories in the usual manner. The software has many inputs, including the software engineers that created a set of rules that described the subject matter (e.g. that the team in baseball with the highest number of runs is the winner). But it also employed what they described as *meta-writers* who created a set of templates that a writer would call upon to create a standard sports story. It is notable that they used the term meta-writers, indicating their creative role in contributing to the software. To add an additional layer of complexity, these meta-writers would often draw upon known sports writers for inspiration for these templates. In the case of baseball, they channelled the writing style of Ring Lardner, who had written over a century earlier.

We can see here that the *AI text*—by which I mean not the final sports stories, but the software system used to generate them—has

many authors: the software engineers, the meta-writers and those whose writings inform the system. Such systems, due to their many human contributors, are more similar to other multi-authored texts than to the sole authored work that so often comes to mind when we consider writing. Indeed, the lack of a single, strong, authorial voice in AI systems contributes to the ability of these systems to hide their authored nature.

We are so used to the simplicity of writing—the act of pen on paper—the idea that the Narrative Science’s system was akin to writing seems odd. But let us remember that other, recently developed, media are more heavily technological and require specialists to achieve their ends. Cinema shows parallels, with a film requiring the expertise from the director, camera-operator, and film editor, to name a few. The situation of multiple authorship is not unlike the multiple authors of many Wikipedia articles. For example, one Wikipedia article on the subject of the ‘B-17 Flying Fortress’ was found in one study to have had (at that time) 376 users who had edited the article, most of these would be minor edits, but a handful of users made considerable input to the text (Jones 2008).

More recently, Elkins and Chun (2020) have explored the potential of the very latest in the AI-based language tools, GPT-3, to generate written material. This technology builds a vast language model based on a corpus of many written texts, drawn from the Internet which themselves previously created by humans, including texts garnered by crawling the web and made available at Common Crawl.⁴ GPT-3 takes the statistical properties of these input texts and employs them for a range of tasks, such as question answering or text generation. The user inputs a starting piece of text and GPT-3 will use its model to output a series of texts that statistically follow-on based on the writings found in the corpus. In this case, instead of leading us to write using conventionally accepted spellings and grammar (found in most text editors), the GPT-3 model will push us towards conventionally accepted lines of thought, and ways of expressing them, that follow from our input. In other words, GPT-3 is a tool to help writing (Floridi and Cjiriati 2020).

⁴ <https://commoncrawl.org/the-data/>.

If GPT-3 is itself a *text*, where are its authors? In this case there are many thousands, if not millions—all those people who own texts were part of the training corpus, as well as the computer scientists that converted these billions of words into the language model. But let us not forget that the language model will contain many of the historical biases and errors that the human authors brought to their original words. It is to these people that we must look to address many of the problems that systems such as GPT-3 have been shown to exhibit. Here is an example where the lessons from the study of existing media can help us understand the smart media. Work on bias in newspapers and television reporting has much in common with the problems now bedevilling AI and machine learning (ML).

AI as a Medium in Images

So far, I have focussed on forms of AI that deal with natural language, partly due to the role of the written word in the Turing Test. However, AI has also had great success in working with images. How does our proposal work in this domain? Let us once again look at some AI systems that work with images and see if we can locate the authors.

One of the most successful ML (a subset of artificial intelligence in which the system ‘learns’ from data presented to it⁵) applications has been in image recognition and classification. For example, AI has been used extensively to help astronomers categorise the types of galaxies seen in telescope images. Galaxies come in a range of morphologies, which reflect their formation history. We all know the classical grand spiral form—such as that of the nearby Andromeda galaxy—but there are many so-called elliptical galaxies which are simpler, ellipsoidal-shaped collections of billions of stars. Both spirals and elliptical galaxies also have properties that lead to more detailed classifications. These classifications are crucial to astronomy, and catalogues of these have been created at

⁵ Notice how even I find it sometimes useful to anthropomorphise AI systems. But I do so ‘knowingly’. See an interesting discussion on anthropomorphising in science by Yanai and Lercher (2020).

great expense of time. Machine learning allows a computer system to take the classifications created by humans and create a model that can then be used to classify any new galaxy image as either a spiral or elliptical type. To do this, humans must provide labels to images in the training set of images, and these data (the images and their labels) are used to generate a model. The conventional position of AI would have it that the computer system is now working as an expert in galaxy classification, as if we have a robot astronomer working next to us. One can almost imagine a press release with cartoon of a robot astronomer looking into a telescope eyepiece.

However, one may ask how such an AI system reflects our notion that AI can be seen as an authored text. Normally many minds are involved in creating the data used in ML (with GPT-3 being the extreme), and normally there is no indication of who these people are. However, for one ML project we find named authors in a similar galaxy classification project, which was used to create a ML model (Baillard et al. 2011). They were: S. Arnouts, A. Baillard, E. Bertin, P. Fouqué, V. de Lapparent, J.-F. Le Borgne, D. Makarov, L. Makarova, H. J. McCracken, Y. Mellier and R. Pelló. In this example, the additional authorial work in creating and refining the machine learning model from the labelled data can be assigned to one of these astronomers, namely Anthony Baillard, as part of his Ph.D. thesis.

If I were to use this ML model to classify my own set of galaxy images it is, I would argue, not akin to working with a robot, but would feel like having the knowledge and experience of these astronomers close to hand. It is as if they had written a handbook on galaxy classification, but the words of the written manual become active in the computer and represent their thinking in a way that is difficult, if not impossible, on paper.

We can go further with authorship—the galaxy classifier built by Dieleman et al. (2015) used data collected by crowdsourcing over half a million Internet users who had used the Galaxy Zoo website⁶ to manually classify galaxies. This collective authorship is similar to the original writers of the texts that made the training corpus for GPT-3 and equally

⁶ <https://www.zooniverse.org/projects/zookeeper/galaxy-zoo/classify>.

makes it hard to see the authors—but they are there, hidden in the weights used in the ML model—invisible to the eye until invoked by running the AI model, much as human voices in a choir can lie silent on a music compact disc until brought to life by a series of complex technological interventions.

Conclusion

This essay argues that AI almost wilfully denies its status as a medium. The creators and users of AI systems play a game in which the users are able to deny the fictional nature of the interaction. Why does it matter whether we see AI as an intelligent agent, rather than a medium? Why not play the game? Indeed, this form of engagement with a technology is one that is found, for example, in viewers of television operas who can treat the world of the TV series *Dallas* as if it were real—up to a point (Ang 1985). We can engage with AI in this spirit, one can imagine an AI form of fiction which deliberately plays with the ELIZA effect. I myself use anthropomorphic language in my own thinking as often it is an easy way to proceed. But we need to do so knowingly.

In the past the denial of the mediating nature of AI systems may have been a small concern. But the current position is different from the early history of AI. It has become far more pervasive: reaching into every area of life. David Watson (2019: 418) puts it well: “while the connections between machine learning algorithms and human cognition may be intriguing and suggestive, the rhetoric of anthropomorphism can do more harm than good when it comes to conceptualizing the important ethical challenges posed by emerging technologies”. I am not suggesting that those who so desire should not continue to strive towards the goal of artificial general intelligence. It is a fascinating intellectual challenge and will doubtless throw light on how we think. However, an issue with the traditional view of AI is that “we also get distracted from questions concerning AI that are relevant right now. These questions have nothing to do with humanoid robots or conscious machines, but with the implementations of ethical values such as fairness, accountability, privacy and transparency” (Hermann 2020: 654).

Taking a new perspective AI can draw upon existing work in the established media: print, radio, cinema and television. To do so will throw light on many of the more urgent issues arising from the widespread use of AI, those questions of fairness, accountability and bias. It will also point the way forward as to how AI systems can be designed to highlight their nature as media. What interfaces are appropriate, for instance?

But more than this, I suggest, a *smart medium* (to coin a phrase) approach to the techniques developed by AI promise so much more—giving humans another, powerful and expressive means of communication. Just as the development of the written word and printing led to literature, and image recording on celluloid led to cinema, how exciting would it be to be involved in the development of the next level of communication media which can enhance our expressive and communicative powers. All too often discussion of AI has focussed on the dehumanising aspects, in which jobs are at risk, and the machines may replace us, a claim Sinclair analyses in this anthology. But we can follow an alternative path.

Bibliography

- Airenti, Gabriella. "The Cognitive Bases of Anthropomorphism: From Relatedness to Empathy." *International Journal of Social Robotics* 7 (2015): 117–127. <https://doi.org/10.1007/s12369-014-0263-x>.
- Ang, Ien. *Watching Dallas: Soap Opera and the Melodramatic Imagination*. Taylor & Francis Ltd., 1985.
- Baillard, Anthony et al. "The FIGI Catalogue of 4458 Nearby Galaxies with Detailed Morphology." *Astronomy and Astrophysics* A74 (2011). <https://doi.org/10.1051/0004-6361/201016423>.
- Bender, Emily M., and Alexander Koller. "Climbing Towards NLU: On Meaning, Form, and Understanding in the Age of Data." In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (July 2020): 5185–5198. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.463>.
- Bolter, Jay D. *Writing Space: The Computer, Hypertext, and the History of Writing*. Lawrence Erlbaum Associates, Inc., 1991.

- Dieleman, Sander, Kyle Willett, and Joni Dambre. "Rotation-Invariant Convolutional Neural Networks for Galaxy Morphology Prediction." *Monthly Notices of the Royal Astronomical Society* 450 (2015): 1441–1459. <https://doi.org/10.1093/mnras/stv632>.
- Elkins, Katherine, and Jon Chun. "Can GPT-3 Pass a Writer's Turing Test?" *Journal of Cultural Analytics* (September 14, 2020). <https://doi.org/10.22148/001c.17212>.
- Floridi, Luciano, and Massimo Chiriatti. "GPT-3: Its Nature, Scope, Limits, and Consequences." *Minds & Machines* 30 (2020): 681–694. <https://doi.org/10.1007/s11023-020-09548-1>.
- Heider, Fritz, and Marianne Simmel. "An Experimental Study of Apparent Behavior." *The American Journal of Psychology* 57 (1944): 243–259. <https://doi.org/10.2307/1416950>.
- Hendler, James. "Avoiding Another AI Winter." *IEEE Intelligent Systems* 23, 2 (March–April 2008): 2–4. <https://doi.org/10.1109/MIS.2008.20>.
- Hermann, Isabella. "Beware of Fictional AI Narratives." *Nature Machine Intelligence* 2 (2020): 654. <https://doi.org/10.1038/s42256-020-00256-0>.
- Huxor, Avon P. *Artificial Intelligence as Writing: Knowledge-Based Hypertext Systems as a Medium for Communication*. PhD thesis, Middlesex University (1994). <https://eprints.mdx.ac.uk/13407/>.
- Jones, John. "Patterns of Revision in Online Writing: A Study of Wikipedia's Featured Articles." *Written Communication* 25, 2 (2008): 262–289. <https://doi.org/10.1177/0741088307312940>.
- Kastner, John, and S. J. Hong. "A Review of Expert Systems." *European Journal of Operational Research* 18, 3 (1984): 285–292. [https://doi.org/10.1016/0377-2217\(84\)90150-4](https://doi.org/10.1016/0377-2217(84)90150-4).
- Levy, Stephen. "Can an Algorithm Write a Better News Story Than a Human Reporter?" *WIRED* (April 24, 2012). <https://www.wired.com/2012/04/can-an-algorithm-write-a-better-news-story-than-a-human-reporter/>.
- Lipton, Zachary, and Jacob Steinhardt. "Troubling Trends in Machine Learning Scholarship: Some ML Papers Suffer from Flaws That Could Mislead the Public and Stymie Future Research." *Queue* 17, 1 (January–February 2019): 45–77. <https://doi.org/10.1145/3317287.3328534>.
- Martin, Nicole. "Did A Robot Write This? How AI Is Impacting Journalism." *Forbes*, February 8, 2019. <https://www.forbes.com/sites/nicolemartin1/2019/02/08/did-a-robot-write-this-how-ai-is-impacting-journalism/>.

- McCarthy, John, M. L. Minsky, N. Rochester, and C. E. Shannon. 2006. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence." *AI Magazine* 27, 4 (August 31, 1955): 12. <https://doi.org/10.1609/aimag.v27i4.1904>.
- McCorduck, Pamela, M. Minsky, O. Selfridge, and H. A. Simon. "History of Artificial Intelligence." In *Proceedings of the 5th International Joint Conference on Artificial intelligence - Volume 2* (IJCAI'77). Morgan Kaufmann Publishers Inc., San Francisco, CA (1977): 951–954.
- McDermott, Drew. "Artificial Intelligence Meets Natural Stupidity." *SIGART Bulletin* 57 (April 1976): 4–9. <https://doi.org/10.1145/1045339.1045340>.
- Moor, James. 2006. "The Dartmouth College Artificial Intelligence Conference: The Next Fifty Years." *AI Magazine* 27, 4: 87–91. <https://doi.org/10.1609/aimag.v27i4.1911>.
- Olson, David R. *The World on Paper: The Conceptual and Cognitive Implications of Writing and Reading*. Cambridge: Cambridge University Press, 1994.
- Paulson, William. "Computers, Minds, and Texts: Preliminary Reflections." *New Literary History* 20, 2 (1989): 291–303. <https://doi.org/10.2307/469102>.
- Peterson, James L. 1980. "Computer Programs for Detecting and Correcting Spelling Errors." *Communications of the ACM* 23, 12 (December 1980): 676–687. <https://doi.org/10.1145/359038.359041>.
- Purtill, Richard L. "Beating the Imitation Game." *Mind*, New Series 80, 318 (1971): 290–294. <http://www.jstor.org/stable/2252425>.
- Reeves, Byron, and Clifford I. Nass. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge: Cambridge University Press, 1996.
- Saygin, A.P., I. Cicekli, and V. Akman. "Turing Test: 50 Years Later." *Minds and Machines* 10 (2000): 463–518. <https://doi.org/10.1023/A:1011288000451>.
- Sokolowski, Robert. "Natural and Artificial Intelligence." *Daedalus* 117, 1 (Winter 1988): 45–64. https://www.jstor.org/stable/20025138?seq=1#metadata_info_tab_contents
- Turing, Alan M. "I.—Computing Machinery and Intelligence." *Mind* 59, 236 (October 1950): 433–460. <https://doi.org/10.1093/mind/LIX.236.433>.
- Watson, David. "The Rhetoric and Reality of Anthropomorphism in Artificial Intelligence." *Minds & Machines* 29 (2019): 417–440. <https://doi.org/10.1007/s11023-019-09506-6>.

- Weizenbaum, Joseph. "ELIZA—A Computer Program for the Study of Natural Language Communication Between Man and Machine." *Communications of the ACM* 9, 1 (January 1966). <https://doi.org/10.1145/365153.365168>.
- Yanai, Itai, and Martin Lercher. "The Two Languages of Science." *Genome Biology* 21, 147 (2020). <https://doi.org/10.1186/s13059-020-02057-5>.
- Złotowski, J., H. Sumioka, and F. Eyssel, et al. "Model of Dual Anthropomorphism: The Relationship Between the Media Equation Effect and Implicit Anthropomorphism." *International Journal of Social Robotics* 10, 701–714 (2018). <https://doi.org/10.1007/s12369-018-0476-5>